

UC Berkeley

UC Berkeley Previously Published Works

Title

Relatedness, relevance and responsiveness in retrieval systems

Permalink

<https://escholarship.org/uc/item/4nh2358q>

Journal

Information Processing & Management, 19(4)

ISSN

0020-0271

Author

Buckland, Michael K

Publication Date

1983

DOI

10.1016/0306-4573(83)90020-1

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial-NoDerivatives License, available at

<https://creativecommons.org/licenses/by-nc-nd/4.0/>

Peer reviewed

## RELATEDNESS, RELEVANCE AND RESPONSIVENESS IN RETRIEVAL SYSTEMS

MICHAEL K. BUCKLAND

School of Library and Information Studies, University of California, Berkeley, CA94720, U.S.A.

(Received for publication 13 December 1982)

**Abstract**—The varying uses of the term “relevance” are considered by analyzing retrieval-related activities into three different processes (Formulation, Retrieval, and Utilization) and four sorts of entities (Enquiry, Attribute, Response, and Benefit).

Three quite different sorts of “relevance” are identified (Responsiveness, Pertinence, and Beneficiality). It is argued that relevance in the sense of utility cannot properly be used to evaluate the performance of retrieval processes as retrieval processes. Any use of utility implies that the utilization is also being considered.

### PROBLEM

The design, use, and evaluation of retrieval systems all depend heavily on various sorts of relatedness. In particular the term “relevance” continues to be freely used for more than one sort of relatedness, including some for which there appears to be no other generally accepted term [1]. One consequence of this semantic problem is that the evaluation of retrieval systems is hindered. An attempt will be made to clarify the concepts and terminology involved.

### DEFINITIONS

We characterize an information retrieval system as a device whereby a set of data is yielded in response to an enquiry. It can happen that the response yields no data—a null set.

The distinction between data retrieval and “document” retrieval is commonly made and three categories can be distinguished:

(i) Data retrieval, whereby signals or data are received with the expectation that they will be informative for whoever posed the enquiry;

(ii) True document retrieval (alias “full text retrieval”), whereby extended texts are retrieved. This is one form of data retrieval;

(iii) Reference retrieval (usually incorrectly referred to as “document retrieval”) whereby the data retrieved can (also) be used as a point to further sources of data, usually external to the retrieval system. For example, if bibliographical data were retrieved, one might properly regard this process as reference retrieval if one’s purpose were to proceed on to the eventual retrieval of the documents cited, but as data retrieval if one were interested in the bibliographical citations themselves rather than the documents cited.

Although these distinctions are important at a practical level in terms of the choice of techniques for implementation, the distinctions appear to be substantially irrelevant to the present discussion.

The mechanism by which retrieval systems operate is the association, usually (but not necessarily) the matching of arbitrarily chosen but predetermined attributes of the set of data that is to be susceptible to retrieval. The attributes that are used include authorship (as in a library catalog), date of publication (as, sometimes, in bibliographies), age (as in museum documentation), occurrence of words (in the searching of texts) and so on. The list of possible attributes that *could* be used seems endless: location, size, chemical process, origin, etc. Nor need attributes be used alone: systems retrieving bibliographical data, for example, commonly operate on two or more attributes in combination; e.g. authorship, date of publication, language, and subject matter. The retrieval process responds to an enquiry by yielding such data as it finds that are highly associated with the attributes specified in the enquiry.

Writings about retrieval and especially about the evaluation of information retrieval systems have been dominated by just one of the apparently unlimited range of attributes: subject matter, i.e. what documents are "about". It is suggested that this preoccupation with one particular case, perhaps because it is technically interesting, has hindered clarity of thought about the foundations of information retrieval theory. However, retrieval by 'aboutness' has to be seen as the use of one attribute among many.

#### THE RETRIEVAL PROCESS

The use of the retrieval system ordinarily includes three distinguishable stages:

(A) *Formulation* of an enquiry in the terms of the retrieval system, i.e. in terms of the attribute(s) used as a basis for retrieval. This can be a non-trivial problem in a system that uses terminology that is linguistically and/or conceptually unfamiliar. The process is sometimes referred to as "translating the enquiry into the language of the indexing system".

(B) *Retrieval* of data by the retrieval system. Data are presented on a screen, found on a card, located on a printed page, and so on. The essential feature is that the system yields a set of data in response to the enquiry.

(C) *Utilization*. The individual who posed the enquiry may become informed as a result of discovering the data yielded and, one hopes, will derive benefit from doing so.

The effectiveness of each process can vary:

The *formulation* of the enquiry may be more or less appropriate to the choice of attributes available in the retrieval system. Much depends on the extent to which the person doing the formulation has an understanding of the data sought and of the characteristics of the retrieval system. The phrasing of the enquiry may be unskillful.

The *retrieval* process may be ineffective in two ways: It may respond with data that have been incorrectly included; and it may fail to yield data that has been wrongly excluded. For example, using authorship as the attribute for retrieval "Mark Twain" might be used as the enquiry. The retrieval system might correctly respond with data concerning works written by Mark Twain but not retrieve data concerning works written by Samuel Clemens—or, if it did, it might also wrongly include data concerning works by Severus Clemens or Susie Clemens.

*Utilization* (not retrieval) might be impaired if the enquirer did not have the necessary abilities (e.g. knowledge, language competence, cognitive skills) to become informed by the data retrieved.

From the analysis thus far there are numerous possibilities for things to be related to one another—or to have degrees of relatedness. These include but are not limited to relationships between: (a) the enquiry as formulated for the retrieval system; (b) any of a seemingly unlimited range of attributes; (c) data retrieved; and (d) benefit to the user.

Before considering sorts of relatedness, it is important to emphasize the separateness of the Retrieval process (B) from the process of Formulation (A) and from Utilization (C). This is reflected in the distinction between the Data retrieved (c) and Benefit to the user (d). The difference can be conveniently illustrated by what we might call the paradox of the disappearing user. Let us imagine that someone posed an enquiry concerning, say, chocolate, cholesterol, and heart disease to a computer-based retrieval system. The retrieval system responds by yielding a set of data. The user becomes better informed as a result of perusing the data and benefits from a changed state of knowledge. Let us now imagine instead that, having posed the enquiry, the enquirer loses interest, is unable to await the response, or dies of a heart attack. In this latter scenario there is no opportunity for utilization (C), nor, therefore, for benefit to the user (d). Yet the retrieval system has performed in exactly the same way. The process of Retrieval (B) and the Data retrieved (c) are indistinguishable, in fact, unchanged from one scenario to the other.

This example illustrates a simple case. In other cases the user may modify the formulation of the enquiry if it is thought that the data yielded would not be what was desired. In this case what has happened is that a *different* search has been formulated however slight the modification has been. Commonly the user's knowledge has changed as a result of preliminary indications concerning the set of data that would be yielded. The

response by the retrieval system to any formulated search will not have been changed unless the retrieval system itself has been altered, e.g. new data added or the indexing modified.

#### RELATEDNESS

In the evaluation of information retrieval systems the term "relevance" has been loosely used to denote differing forms of relatedness. It is suggested that a useful approach would be to define the most useful relationships and degrees of relatedness first and then give them distinctive names afterwards. We shall consider three different relationships for each of which the terms 'relevant' or 'relevance' has been used.

##### (i) *Relatedness I: responsiveness*

Firstly we consider the relatedness of the data retrieved to the enquiry as posed in terms of the attributes used as the basis for retrieval. To what extent *did* the system retrieve all and only the works by Mark Twain? The quality of the response by the retrieval system to the query will be affected by several factors including the appropriateness and completeness of the database, the suitability of the attribute(s) used as the basis for retrieval (in this case authorship), and the ability of the retrieval system to identify those data that fit the description offered by the enquiry—or fit it to the desired degree. If one wished to avoid talking of the 'relevance' of the retrieval data to the enquiry, one might speak of the *responsiveness* of the system. Some term is needed for this sort of relatedness.

##### (ii) *Relatedness II: pertinence*

In (i) above, we were concerned with a general term: the relatedness of the enquiry (a) to the system's response (c) regardless of the attribute(s) being used as a basis for retrieval. We now consider one special case within the general class: when the attribute used as the basis for retrieval is topicality—the subject matter of the data, what the data are "about". In ordinary speech one might well speak of an article as being "relevant" to a given topic. Such relationships are not always straightforward (e.g. general to specific, overlap) and can be difficult to understand or to define. In terms of the present discussion this sort of incomplete relatedness would seem best described as the relatedness between the properties of data with respect to the attributes used for retrieval when these properties are related but not identical. For example, when retrieving by the attribute of topicality, data on "Freud" is related to ("relevant to") the topic of psychoanalysis but is not identical. This kind of relatedness between properties of data with respect to a given attribute could in principle exist with other attributes than topicality. We might term this relatedness "pertinence".

##### (iii) *Relatedness III: beneficiality*

A relationship that is entirely different from either of the above is that between the retrieved data (c) and the benefit of the user (d). It is in this sense, for example, that Wilson has sought to limit the use of the term "relevance" in *Two Kinds of Power* [2] and this sense is implicit in all discussions of "utility-theoretic indexing". Stated simply, it is assumed that retrieval systems are provided and used in order that their utilization will have beneficial effects. Social values are implied. This raises two questions: (1) Whose values? The users' or those of the providers of the service? (Information retrieval theorists tend to assume that the user is—or should be—the arbiter of utility, but this is open to argument); and (2) are we referring to actual benefits or expected benefits? These are critical questions. However, whatever answers are given it is clear that the relationship is different in kind from both (i) and (ii) above because factors external to the retrieval system affect the outcome: Social values; and the knowledge and cognitive skills of the users [3].

#### CONCLUSIONS

The foregoing discussion leads to the following conclusions:

(1) The three sorts of relatedness: responsiveness, pertinence, and beneficiality—are different in kind:

(2) Since beneficiality (Relatedness III) is rooted in the utilization (C) of retrieved data (c) and in human values, the superficially startling conclusion follows that *relevance in the sense of utility cannot properly be used to evaluate the performance of retrieval processes as retrieval processes*. In this we are using a strict, narrow definition: the ability of the system to yield a set of data responsive to the formulated enquiry. The proper basis for evaluation would be rooted in the responsiveness of the system in terms of the enquiry and the attributes used: "fitting the description". Any use of beneficial effects (only) cannot be an evaluation of retrieval processes but must be either an evaluation of utilization of retrieved data or some combination of utilization and other processes, e.g., retrieval (B) and utilization (C) and, possibly, formulation (A)—and even other logically prior activities (e.g. identification of need) or concurrent activities (e.g. improvement of cognitive skills in order to improve utilization) as well [4]. Whoever uses benefit or utility should specify the boundaries of what is being evaluated.

(3) New distinct terms should be adopted to distinguish different sorts of relatedness ("relevance") if confusion is to be reduced. Ideally the terms should be applicable generally to sorts of relatedness and not restricted to specific examples of retrieval activities (e.g. subject-based retrieval as opposed to retrieval on the basis of other attributes) unless such narrower usage is necessary or justified and clearly indicated.

#### DISCUSSION

The distinction made above between Relatedness I: Responsiveness and Relatedness III: Beneficiality resembles Orr's distinction between two kinds of "library goodness" [5]:

(1) How good is it? A measure of quality and, therefore, of capability.

(2) What good does it do? A measure of value and, therefore, of beneficial effects.

The latter is unavoidably rooted in social values of some kind; the former is value-neutral in terms of human values as opposed to techniques.

It does not follow from any of this that retrieval based on responsiveness and pertinence cannot be related to beneficial effects. Value-neutral attributes can be used as predictors of probable utility. Retrieval of data on busses on the attribute of "time since last servicing" is a good but not perfect predictor of which vehicles should be serviced next in order to minimize maintenance costs. A combination of vehicle model, age, and time since last service might permit even more useful predictions. None of this affects the accuracy of the retrieval, although the accuracy of the retrieval (responsiveness) will affect accuracy and, therefore, the benefit derived from the predictions. Similarly, it is a reasonable prediction that someone seeking to learn about psychoanalysis would find a book about Freud useful. It does not follow that the book about Freud should have been indexed as if about psychoanalysis. Consistency in indexing by attributes rather than by imagined future utility enables the *user* rather than the indexer to use attributes as a basis for deriving from the retrieval system what is likely to be useful. In this way retrieval systems that are consistent and reliable in their responsiveness in terms of whatever attributes are used form a good basis for the retrieval of data that will, in turn, form a good basis for users to derive benefit.

#### REFERENCES AND NOTES

- [1] For a convenient introduction to the literature concerning relevance see A. BOOKSTEIN, *Relevance. J. Am. Soc. Inform. Sci.* 1979 30(5), 269–273.
- [2] P. WILSON, *Two Kinds of Power: An Essay on Bibliographical Control*. University of California Press (1968). See especially chapters II Describing and Exploring; and III Relevance.
- [3] In this discussion we are concerned with the knowledge and cognitive skills of the users as users. This is separable in theory from the skills of whoever formulates the enquiry even though both roles could be performed by the same individual.
- [4] The difficulties encountered in trying to use beneficiality where responsiveness rather than beneficiality is appropriate is illustrated in a recent discussion of the retrieval of data that are relevant (i.e. beneficial) but not topical. "The material is relevant but not topical. However, to expect a retrieval system to respond to an unexpressed need seems a harsh requirement indeed . . . the instance of the relevant but untopical document is unlikely to occur in any test of relevant

documents used for system evaluation. . . while it is certainly the case that topicality is not a necessary condition for relevance, it seems that we comfortably treat it as such without great loss.”  
B. BOYCE, Beyond Topicality; a two stage view of relevance and the retrieval process. *Inform. Proc. Management* 1982, **18**(3), 105–109.

- [5] R. M. ORR, Measuring the goodness of library services: a general framework for considering quantitative measures. *J. Documentation* 1973, **29**(3), 315–332.